


Start where you are — four stages, not a big bang

Each stage moves you one maturity level: see it, bound it, integrate it, prove it.



START MONDAY

Pick one real AI use case › classify & risk-tier it › select the controls for that tier › collect evidence & grade the boundary › prove it, then repeat.

Where is your organisation today?

Read each level. You are at the highest level where the statements are true.

1	Ad Hoc	☐ AI use isn't consistently inventoried; ownership is unclear.	☐ Vendor AI features get enabled by default.	☐ No one can reconstruct what AI saw, produced, or triggered.
2	Visible	☐ An AI inventory exists; high-risk use cases are identified.	☐ A risk-tiering model exists but isn't applied consistently.	☐ Evidence exists but is scattered; shadow AI remains.
3	Integrated	☐ Risk tier drives the required controls.	☐ AI review is part of architecture, IAM, data governance, vendor risk.	☐ Exceptions are time-bound and owned; high-risk AI has evidence.
4	Measurable	☐ Control coverage is measured by pillar and tier.	☐ Assurance testing produces repeatable evidence; audit can review it.	☐ Kill-switches and rollback are tested where required.
5	Adaptive	☐ Monitoring signals and incidents update control priorities.	☐ Vendor or model changes trigger review automatically.	☐ Controls continuously improve from findings and near-misses.